

AT A GLANCE

INDUSTRY

Aerospace & Defense

PRIMARY WORKLOAD

Autonomous agent & engineering tooling

MODELS IN PRODUCTION

GPT-5.4 · GPT-5.1 · GPT-4.1

COMPLIANCE BOUNDARY

ITAR

KEY RESULTS

1

API for every frontier model they use

0

Prompts used to train a model

<24 HRS

From feature request to shipped fix

IN THEIR WORDS

“

LLMs have been a big accelerator for developing our cutting-edge re-entry vehicles, and Consus has been the accelerator we need to get the latest models as quickly as possible while still maintaining compliance.

Alex Tacescu, Director of Software

Inversion Space

CASE STUDY · AEROSPACE & DEFENSE

EVERY FRONTIER MODEL ONE COMPLIANT API.

Inversion Space is building vehicles that return cargo from orbit to anywhere on Earth in under an hour. To keep that pace, their engineers reach for frontier AI every day. The last thing they wanted was to become AI-infrastructure and compliance experts along the way, or to put proprietary work at risk to move fast.

01 · THE CHALLENGE

A NEW MODEL SHOULDN'T MEAN A NEW PROJECT

The best model for a job changes every few weeks, and Inversion's engineers wanted to reach for whichever one fit, whether the work was autonomy, simulation, or everyday software. But each provider brought its own integration to build, its own credentials to hold, and its own compliance question to answer. Those questions only got sharper with proprietary vehicle and mission data in play. Keeping up with the frontier was turning into a stream of infrastructure and compliance work that had nothing to do with flying hardware. That was never the work Inversion set out to do.

02 · THE CONSUS APPROACH

ONE API, INSIDE A BOUNDARY THEY CAN TRUST

Consus gives Inversion's engineers a single API that reaches compliant infrastructure using a standard request format, so trying a different model is a one-line change instead of a new integration. Inversion runs the GPT family in production, including **GPT-5.4** and **GPT-5.1**, to meet both performance and compliance requirements. Every request runs inside ITAR-eligible environments operated entirely by U.S. persons, and nothing they send is used to train a model. Their proprietary work can reach frontier AI without a fresh security review each time the team wants to try something new.

03 · WORKING TOGETHER

A TEAM ON THE OTHER END, NOT A TICKET QUEUE

The part that is hardest to replace is the relationship. When Inversion's engineers needed more room for longer autonomous runs, Consus turned the fix around in under a day. New models show up as soon as they clear the compliance bar. When a question comes up, Inversion talks to the people who build the gateway, not a support line.